

IN THE CLAIMS:

1 1. (Currently Amended) A storage system for use in a storage system cluster, the storage
2 system comprising:

3 a first server and a second server, wherein the second server is a cluster partner to
4 the first server; and

5 a storage operating system operating on the first server, the storage operating sys-
6 tem including a cluster connection manager adapted-configured to create, destroy, and
7 maintain one or more communication sessions with a-the cluster partner, the cluster con-
8 nection manager operatively interconnected with a set of cluster connection manager cli-
9 ents, where each cluster connection manager client is a process executing on the storage
10 system, and wherein the cluster connection manger is further configured to create, de-
11 stroy, and maintain a virtual interface connections between a cluster connection manager
12 client on the first server with a cluster connection manager client on the second server to
13 form a peer process between the cluster connection manager clients.

1 2. (Original) The storage system of claim 1 wherein one of the set of communication
2 clients comprises a failover monitor.

1 3. (Original) The storage system of claim 1 wherein one of the set of cluster connection
2 manager clients comprises a non-volatile random access memory shadowing process.

1 4. (Original) The storage system of claim 1 wherein the cluster connection manager is
2 further adapted to perform connection management operations in response to communi-
3 cations from the connection manager clients.

1 5. (Original) The storage system of claim 4 wherein the communications comprise an ap-
2 plication program interface function call.

1 6. (Original) The storage system of claim 1 wherein the cluster connection manager is
2 further adapted to load balance the one or more communication sessions over a plurality
3 of cluster interconnect devices.

1 7. (Original) The storage system of claim 1 wherein the cluster connection manager is
2 further adapted to perform a failover procedure for one or more communication sessions
3 from a failed cluster interconnect device to an operational cluster interconnect device.

1 8. (Original) The storage system of claim 1 wherein the cluster connection manager is
2 operatively interconnected with a plurality of cluster interconnect devices.

1 9. (Original) The storage system of claim 1 wherein the storage operating system com-
2 prises a plurality of cluster connection managers.

1 10. (Currently Amended) A storage operating system, executing on a storage system, the
2 storage operating system comprising:

3 | a cluster connection manager ~~adapted~~ configured to manage a set of peer-to-peer
4 | connections associated with a set of cluster connection manager clients executing on the
5 | storage system, wherein the cluster connection manger is further configured to create,
6 | destroy, and maintain a virtual interface connection between a cluster connection man-

7 | ager client on a first server with a cluster connection manager client on a second server to
8 | form a peer process between cluster connection manager clients.

1 | 11. (Original) The storage operating system of claim 10 wherein the set of cluster con-
2 | nection manager clients comprises a failover monitor.

1 | 12. (Original) The storage operating system of claim 10 wherein the cluster connection
2 | manager is further adapted to perform load balancing of the set of peer-to-peer connec-
3 | tions over a plurality of cluster interconnect devices.

1 | 13. (Original) The storage operating system of claim 10 wherein the cluster connection
2 | manager is further adapted to failover the set of peer-to-peer connections from a failed
3 | cluster interconnect device to an operational cluster interconnect device.

1 | 14. (Currently Amended) A method for initiating a peer-to-peer communication session,
2 | ~~the method comprising the steps of:~~

3 | creating, using a cluster connection manager executing on a ~~storage system~~first
4 | server, an initial connection with a cluster partner on a second server;
5 | exchanging a set of peer connection information;
6 | passing a set of cluster connection manager client information to the cluster part-
7 | ner, wherein the set of cluster connection manager client information includes at
8 | least one virtual interface and any memory requirements for each cluster manager
9 | client;
10 | creating a set of appropriate communication ports using the set of cluster connec-
11 | tion manager client information;
12 | alerting the cluster partner of a ready status; and

13 | alerting a set of cluster connection manager clients that the cluster partner is in a
14 | ready state.

1 | 15. (Original) The method of claim 14 wherein the set of clients comprises a failover
2 | monitor process.

1 | 16. (Original) The method of claim 14 wherein the set of peer connection information
2 | comprises a version number.

1 | 17. (Currently Amended) The method of claim 14 wherein the step of passing a set of
2 | client information to the cluster partner further comprises ~~the steps of:~~
3 | collecting, from a set of clients, the set of client information; and
4 | transferring the collected set of client information to the cluster.

1 | 18. (Original) The method of claim 17 wherein the client information comprises a num-
2 | ber of communication ports required.

1 | 19. (Original) The method of claim 17 wherein the set of client information further com-
2 | prises an amount of memory requested by a particular client.

1 | 20. (Currently Amended) The method of claim 14 wherein the step of creating an initial
2 | connection further comprises ~~the step of~~ using remote direct memory access primitives to
3 | create the initial connection.

1 21. (Currently Amended) The method of claim 14 wherein the step of creating an initial
2 connection further comprises ~~the step of~~ performing a series of remote direct memory
3 access operations to create the initial connection.

1 22. (Currently Amended) A method for terminating a peer-to-peer communication ses-
2 sion, ~~the method comprising the steps of:~~
3 alerting, using a cluster connection manager executing on a storage system, a set
4 of clients of an impending termination of the communication session;
5 closing, by the clients, a set of communication ports associated with the commu-
6 nication session, wherein the set of communication ports comprise a set of virtual inter-
7 face connections; and
8 performing an initialization of a peer-to-peer communication session procedure.

1 23. (Original) The method of claim 22 wherein the set of communication ports comprises
2 a set of virtual interface connections.

1 24. (Original) The method of claim 22 wherein the set of clients comprises a failover
2 monitor.

1 25. (Currently Amended) A storage ~~operating system, executing on a storage system, the~~
2 ~~storage operating system system cluster,~~ comprising:
3 a first storage system of the storage system cluster, the first storage system having
4 a first disk shelf and a first cluster connection manager to manage data flow from/to an
5 external source to the first disk shelf, wherein the first cluster connection manger is con-
6 figured to create, destroy, and maintain a virtual interface connections between a cluster

7 connection manager client on the first storage system with a cluster connection manager
8 client on a second storage system to form a peer process between the cluster connection
9 manager clients;

10 a second storage system of the storage system cluster, the second storage system
11 having a second disk shelf and a second cluster connection manager to manage data flow
12 from/to an external source to the second disk shelf; and

13 the first cluster connection manager to shift data flow from/to the first disk shelf
14 to the second disk shelf upon an event condition. ~~a cluster connection manager having~~
15 ~~means to manage a set of peer-to-peer connections associated with a set of cluster con-~~
16 ~~nection manager clients executing on the storage system.~~

1 26. (Currently Amended) The storage system cluster operating system of claim 25
2 wherein the event condition is a failed interconnect driver connected the first cluster con-
3 nection manager~~the set of cluster connection manager clients further comprises a failover~~
4 ~~monitor.~~

1 27. (Currently Amended) The storage system cluster operating system of claim 25
2 wherein the event condition is a load-balancing condition~~the set of cluster connection~~
3 ~~manager clients further comprises a nonvolatile random access memory shadowing proc-~~
4 ~~ess.~~

1 28. (Currently Amended) A system configured to manage reliable peer communication
2 among storage systems in a clustered environment, the system comprising:

3 one or more peer processes executing on each storage system partner; and

4 a cluster connection manager executing on each storage system partner, the clus-
5 ter connection manager creating a set of peer-to-peer connections between the one or
6 more peer processes executing on each storage system, wherein the cluster connection

7 | manager is provided to reliably create virtual interface connections between peer proc-
8 | esses executing on the storage system partners over a cluster interconnect without requir-
9 | ing a storage operating system executing on each storage system to be fully active or
10 | functioning.

1 | 29. (Currently Amended) A computer readable medium for initiating a peer-to-peer
2 | communication session, the computer readable medium including program instructions
3 | executed by a processor for performing the steps of:
4 | creating, using a cluster connection manager executing on a ~~storage system~~ first
5 | server, an initial connection with a cluster partner on a second server;
6 | exchanging a set of peer connection information;
7 | passing a set of cluster connection manager client information to the cluster part-
8 | ner, wherein the set of cluster connection manager client information includes at
9 | least one virtual interface and any memory requirements for each cluster manager
10 | client;
11 | creating a set of appropriate communication ports using the set of cluster connec-
12 | tion manager client information;
13 | alerting the cluster partner of a ready status; and
14 | alerting a set of cluster connection manager clients that the cluster partner is in a
15 | ready state.

1 | 30. (Currently Amended) A computer readable medium for terminating a peer-to-peer
2 | communication session, the computer readable medium including program instructions
3 | executed by a processor for performing the steps of:
4 | alerting, using a cluster connection manager executing on a storage system, a set
5 | of clients of an impending termination of the communication session;

6 closing, by the clients, a set of communication ports associated with the commu-
7 nication session, wherein the set of communication ports comprise a set of virtual inter-
8 face connections; and
9 performing an initialization of a peer-to-peer communication session procedure.

1 31. (Previously Presented) A method for maintaining a peer-to peer communication, the
2 method comprising:

3 waiting for an event from a client communicating with a cluster partner to be re-
4 ceived by a cluster connection manager executing on a storage operating system;
5 determining whether the event is a client event; and
6 in response to determining that the event is a client event, performing the event
7 utilizing the cluster connection manager.

1 32. (Previously Presented) The method of claim 31, further comprising:

2 in response to determining that the event was not a client event, alerting a set of
3 clients of an impending termination of the communication session;

4 closing, by the clients, a set of communication ports associated with the commu-
5 nication session; and

6 performing an initialization of a peer-to-peer communication session procedure.

1 33. (Previously Presented) The method of claim 32 wherein the set of communication
2 ports comprises a set of virtual interface connections.

1 34. (Previously Presented) The method of claim 32 wherein the set of clients comprises a
2 failover monitor.

1 35. (Previously Presented) The method of claim 31 further comprising monitoring the
2 status of one or more cluster interconnect drivers utilizing the cluster connection man-
3 ager.

1 36. (Currently Amended) A computer readable medium for maintaining a peer-to-peer
2 communication session, the computer readable medium including program instructions
3 | executed by a processor for performing the steps of:

4 waiting for an event from a client involved in a communication session to be re-
5 ceived by a cluster connection manager executing on a storage operating system;
6 determining that the event is a client event; and
7 in response, performing the event utilizing the cluster connection manager.

1 37. (Previously Presented) A storage operating system, executing on a storage system, the
2 storage operating system comprising:

3 one or more peer processes executing on each storage system partner;
4 a plurality of cluster interconnect drivers executing on the storage system; and
5 one or more cluster connection managers configured to detect a failure of a first
6 cluster interconnect driver and in response to determining the failure of the first cluster
7 interconnect driver, utilize a second cluster interconnect driver to access each storage sys-
8 tem partner.

1 38. (Previously Presented) A storage operating system, executing on a storage system, the
2 storage operating system comprising:

3 one or more peer processes executing on each storage system partner;
4 a plurality of cluster interconnect drivers executing on the storage system; and

5 one or more cluster connection managers configured to detect a high bandwidth
6 load on a first cluster connection manager and in response to detecting a high band width
7 load, utilize a second cluster connection manager to access each storage system partner.